1989-03

Conversion of mass storage hierarchy in an IBM computer network

Mauck, Linda S.

Monterey, California. Naval Postgraduate School

http://hdl.handle.net/10945/25792

NAVAL POSTGRADUATE SCHOOL
Monterey, California

THESIS

CONVERSION OF MASS STORAGE HIERARCHY IN AN IBM COMPUTER NETWORK

by

Linda S. Mauck

March 1989

Co-Advisor D. G. Williams
Co-Advisor N.F. Schneidewind

Approved for public release; distribution is unlimited.

REPORT DOCUMENTATION PAGE

1a Report Security Classification Unclassified	1b Restrictive Markings
2a Security Classification Authority	3 Distribution/Availability of Report Approved for public release; distribution is unlimited.
2b Declassification/Downgrading Schedule	
4 Performing Organization Report Number(s)	5 Monitoring Organization Report Number(s)
6a Name of Performing Organization Naval Postgraduate School / 6b Office Symbol *(if applicable)* 54	7a Name of Monitoring Organization Naval Postgraduate School
6c Address *(city, state, and ZIP code)* Monterey, CA 93943-5000	7b Address *(city, state, and ZIP code)* Monterey, CA 93943-5000
8a Name of Funding/Sponsoring Organization / 8b Office Symbol *(if applicable)*	9 Procurement Instrument Identification Number
8c Address *(city, state, and ZIP code)*	10 Source of Funding Numbers Program Element No / Project No / Task No / Work Unit Accession No

11 Title *(include security classification)* CONVERSION OF MASS STORAGE HIERARCHY IN AN IBM COMPUTER NETWORK

12 Personal Author(s) Linda S. Mauck

13a Type of Report	13b Time Covered	14 Date of Report *(year, month, day)*	15 Page Count
Master's Thesis	From To	March 1989	48

16 Supplementary Notation The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

17 Cosati Codes			18 Subject Terms *(continue on reverse if necessary and identify by block number)*
Field	Group	Subgroup	DASD, migration, MSS

19 Abstract *(continue on reverse if necessary and identify by block number)*

With the shift from batch applications to online systems supporting the strategic role of information, corporate or institutional goals tie directly to the information management functions. This has been true at the Naval Postgraduate School (NPS). Like many other Government installations, the NPS Computer Center has to meet its objectives with less than state-of-the-art hardware. In the early 1980's, the Center employed IBM's 3850 Mass Storage Subsystem (MSS) for online storage of student and faculty data sets. It was installed in December 1980 and performed well for over six years. Faced with IBM's announcement (in February 1985) of the limited future connectivity and compatibility and the increasing maintenance costs, the decision was made to replace the MSS with a hardware/software alternative that would use a more modern and reliable architecture. The objective of this thesis is to define the solution, the data set migration process, and describe the early experience with a multi-level, software-managed, storage system.

20 Distribution/Availability of Abstract ☒ unclassified/unlimited ☐ same as report ☐ DTIC users	21 Abstract Security Classification Unclassified	
22a Name of Responsible Individual D. G. Williams	22b Telephone *(include Area code)* (408) 646-2572	22c Office Symbol 54Ss

DD FORM 1473,84 MAR 83 APR edition may be used until exhausted
All other editions are obsolete

Approved for public release; distribution is unlimited.

Conversion of Mass Storage Hierarchy in an IBM Computer Network

by

Linda S. Mauck
Civilian, Naval Postgraduate School
B.S., University of Oklahoma, 1968

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN INFORMATION SYSTEMS

from the

NAVAL POSTGRADUATE SCHOOL
March 1989

ABSTRACT

With the shift from batch applications to online systems supporting the strategic role of information, corporate or institutional goals tie directly to the information management functions. This has been true at the Naval Postgraduate School (NPS). Like many other Government installations, the NPS Computer Center has to meet its objectives with less than state-of-the-art hardware. In the early 1980's, the Center employed IBM's 3850 Mass Storage Subsystem (MSS) for online storage of student and faculty data sets. It was installed in December 1980 and performed well for over six years. Faced with IBM's announcement (in February 1985) of the limited future connectivity and compatibility and the increasing maintenance costs, the decision was made to replace the MSS with a hardware/software alternative that would use a more modern and reliable architecture. The objective of this thesis is to define the solution, the data set migration process, and describe the early experience with a multi-level, software-managed, storage system.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ACKNOWLEDGEMENTS

I wish to express my gratitude and deep appreciation to all those people whose assistance and support have enabled me to complete this thesis and the courses which were required for this degree. Much gratitude goes to my supervisor, David Norman, for his technical support and for insisting that I perform every step that I possibly could in the process of this conversion. In so doing, I learned much more than any regular student could have done. Without his friendship, support, and sharing of expertise, the project would have bogged down at many points. I owe Professor Doug Williams, Director of the W. R. Church Computer Center, much for his perfectionist eye which enabled me to present my work in a more polished and professional manner. The work he did as my advisor was done on his own time, not as part of his job. I owe Ruth Roy a great deal for teaching me how to use SAS for the analyses in the follow-up portion of my project, but more so, I owe her for the moral support needed when I was tired and discouraged. She, more than any other friend, helped keep my eye on my long term goals.

To my children, Kirsten and Andy, I owe much thanks and more for accepting my goals and allowing me to be mother and student, without shortchanging either vocation. And, to my husband, Charlie, I owe the greatest debt, for supporting me and encouraging me in all that I do, to accomplish much more than I thought I could.

I. INTRODUCTION

Data processing has evolved from primarily accounting-oriented applications to the support of integrated information systems. Conversion from batch applications to on-line management information systems directly ties institutional goals to these information management functions. The efficiency of this management directly impacts an institution's success. This has been true at the Naval Postgraduate School (NPS). As with many other Government installations, the NPS Computer Center has to meet its objectives with less than state-of-the-art hardware. The number and size of the data sets belonging to the students, staff, and faculty of NPS, tenant commands, and other users of the NPS Computer Center was a good fit for the IBM 3850 Mass Storage Subsystem (MSS). The MSS was installed in December 1980, between academic quarters, and functioned effectively for six-and-a-half years. IBM announced in February 1985 that no mainframe Central Processing Unit (CPU) beyond the model 3090 would support the MSS. [Ref. 1] This fact plus the increasing maintenance costs ($82,000 for 1988) caused the Center's management to explore alternative hardware/software solutions for a more modern and reliable architecture which promised future connectivity as well. The objective of this thesis is to define the solution and the NPS Computer Center's migration to it. The thesis covers all aspects of the complex process from planning and estimation of storage requirements, data set migration, and post-installation experience with the new system. All steps in the installation process were performed by the author unless otherwise noted.

II. INFORMATION STORAGE—BACKGROUND

In a large system environment, information management depends crucially on cost-effective information storage and retrieval. In the 1980's, with the explosive growth of machine-readable information, various data storage systems have evolved. The current options may be portrayed as a storage hierarchy (Figure 1 on page 3) with each level in the hierarchy having different levels of performance, capacity, and price. [Ref. 2, 3, and 4] Many writers define this hierarchy with greater or fewer levels depending upon the products supported by the writer's company. One author had a bottom layer of printed output for data stored in hard-copy form. As storage devices change, the hierarchy may change in implementation, but these general categories will remain with new levels dependant on cost/performance factors added with technological advancement. The orientation of this thesis is toward the IBM large systems storage hierarchy. Other venders' systems use different approaches.

After processor storage, the top level in today's hierarchy is *high-performance direct-access storage device (DASD)* (solid state) which was first delivered in 1979 to facilitate system paging, a major performance bottleneck in modern systems. Today's online, response-oriented systems require high subsystem availability which can be met by solid-state devices having relatively few mechanical components. Solid-state (non-rotating) DASD is a top performer. Its consistent I/O response time of 0.3 milliseconds (ms) satisfies between 300 and 400 I/O requests per second per I/O path. [Ref. 2, 3, and 4] With the introduction of the 3380 disk storage image in 1984, faster response time for a broad range of online applications became a possibility. According to Mr. Fred Moore of StorageTek, "The provision for device images that mirrored rotating DASD may have been the most significant enhancement for high-performance DASD." The new format allowed the portability of data between real 3380-class DASD and high-performance 3380 without converting blocksizes or changing space allocations via job control language (JCL). Since 1984, solid-state DASD has become the preferred solution for response-critical applications. Although the most costly, the possibilities of 100 percent space utilization and 70 percent channel utilization could make the efficiency of this architecture cost-effective. [Ref. 2]

The second level in the storage pyramid is *cached DASD* controllers whose acceptance has steadily grown throughout the 1980s. Cached controllers serve as high

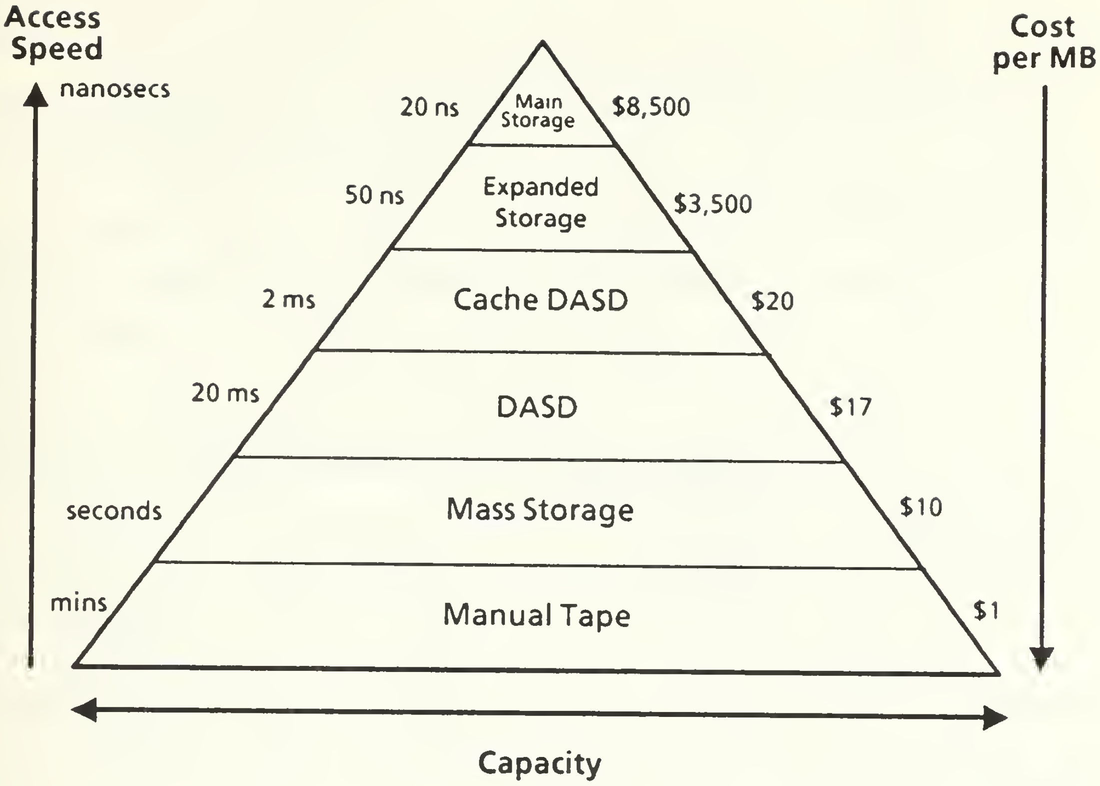

Figure 1. Information Storage Hierarchy

performance holding areas for data that have high READ hit ratios. A WRITE operation must always go to the physical DASD volume for data integrity. A cached subsystem can provide a more consistent response for the attached DASD subsystem (0.3 - 1.0 ms service time) and can improve channel utilization above the typical 35 percent busy threshold for non-cached DASD. With the growth of online and response-critical applications, the use of cache has spread rapidly. Some applications which would be good choices for cached DASD are read-only data sets and databases, database indices, pageable link pack area (PLPA), and catalogs. [Ref. 2] In IBM's hierarchy [Ref. 3] high-performance cached DASD subsystems are represented by IBM 3880 and 3990. The IBM 3880 Model 11 and Model 13 subsystems contain two cached storage directors and a subsystem storage unit with 3350 disks for the Model 11 for paging applications and with 3380 disks for the Model 13 for non-paging applications.

Each storage director attaches to 3-megabytes-per-second, data-streaming channels to attach to DASD. The IBM 3990 Storage Control family replaces the IBM 3880 Storage Control Models 3 and 23. It offers improved price/performance, service, and function over the 3880 and is available in six models: two without cache and four with cache. Five of the models offer four-path access to the new IBM 3380 Enhanced Subsystem DASD. All models attach to the new J and K models as well as older models of IBM 3380 DASD. [Ref. 5]

The third level in the hierarchy is rotating *DASD*, the primary online storage device in almost all computer systems. IBM reports that in 1978, the median number of 3380 disks per IBM installation was approximately nine volumes. This number had grown to over 150 volumes by 1985. [Ref. 6] DASD use has grown at 35 percent annually [Ref. 2] and is predicted to continue its rapid growth at over 30% annually [Ref. 7] or to as much as 45% for some installations [Ref. 6]. DASD satisfies both online and batch requirements with adequate performance capacity and non-volatility. A GUIDE survey published in late 1984 indicated that the amount of allocated space per access mechanism was declining steadily. [Ref. 2] Users have allocated less data per access mechanism to reduce contention and thereby maintain acceptable performance levels. IBM [Ref. 7] predicted that DASD and processor speeds will increase sufficiently to allow the user to allocate 3.5 times the amount of data for comparable response times in the 1990-1995 timeframe.

The fourth level in the hierarchy is *mass storage*. In the late 1960s IBM introduced the 2321 Data Cell for large computer users. Several mass storage products have been introduced since that time though none have yet become dominant for mass storage. At the Naval Postgraduate School, the IBM 3850 Mass Storage Subsystem filled this niche from December 1980 until July 1987. Whereas in the first three levels of the hierarchy, access times are measured in milliseconds, the access time for the mass storage device is measured in seconds. Mass storage subsystems have had problems with reliability and availability to the extent that some companies have discontinued them. Some systems, like IBM's MSS, have used a combination of accessors (picker arms) and data recording devices to transfer data from unique data storage cartridges or high-density video tape stored in a library to staging devices. Other mass storage systems use some industry-standard media, such as 9-track tape reels or the 18-track cartridges, which allow the tapes to be read or written on any compatible drive when there is a subsystem hardware or software failure. In recent years, the successful application of robotics in mass storage subsystems, along with the ability to connect several

subsystems together, gives creators of these systems hope for extensive future use. [Ref. 2] Mass Storage Systems should provide:

- Relatively quick access
- Data access compatible with systems software and access methods
- Technologies which can be enhanced to provide long-term storage solutions
- System accessible storage media
- Cost effectiveness not only in price but operational and environmental measures.

Table 1 summarizes the relationships between these levels in the hierarchy. [Ref. 2, 4, 8, 9, and 10]

Table 1. CHARACTERISTICS OF STORAGE DEVICES

STORAGE DEVICE	CHARACTERISTICS		
	Chan Busy (Percent)	I/O Rate (4Kb/Sec)	Initial Service Time
Solid-state DASD	70-75	750-1500	.3 ms
Cached DASD	50-60	750	.3-1.0 ms
DASD	35-45	375-750	24-33 ms
MSS	80-90	200	2-46 sec

The bottom level of the hierarchy is *magnetic tape*. This removable, portable medium has been the choice for over 20 years for backup, archiving, and transportability. Today, the newer, 18-track cartridge tape subsystems with its 200 megabytes capacity and negligible errors are replacing the traditional 9-track, reel-to-reel tapes which hold 160 megabytes. Mr. Moore forecast [Ref. 2] increasing the 18-track to a 36-track tape and using longer tapes which could increase the capacity of each cartridge to 1.0 gigabytes.

As information becomes more strategic to business, so does the question of recovery from loss of such information. Unlike DASD, tape capacity for business applications is virtually limitless. There will always be a need for this level in the hierarchy and tape, in some form, will be the medium to fit the requirements for many years to come.

If the requirement exists for immediate availability, the data would need to reside on high performance DASD—or the top level of the hierarchy. Level three, DASD, is the choice if a few milliseconds in additional response time can be tolerated. It would

be desirable to have everything instantly accessible but the high cost is not necessary for most data. Mass storage systems satisfy the requirement for lower cost but with an initial service time of several seconds. According to Mr. Moore, "what remains to be seen is a truly successful implementation" of a mass storage subsystem. He also predicts that although the "challenge of managing more than 1,000 gigabytes of online data will be aided to some degree by technology advances, ... the major responsibility will fall heavily into the area of software." [Ref. 2] The reliance only on a hardware hierarchy will cease as software plays an increasingly important part in the future. The requirements of such software will be addressed in Chapter IV.

III. MASS STORAGE SYSTEM AT NPS

When the IBM 3850 Mass Storage Subsystem (MSS), was first marketed by IBM (Oct 9, 1974), the total volume of data collected and maintained by many customers exceeded the maximum configuration of then-current DASD. While the IBM 370/168 processors and 3350 DASD were relatively new, IBM announced the MSS as a tape replacement providing almost unlimited data storage online and at a very low cost [Ref. 11]. At the time, the only alternative was massive off-line tape libraries. There are many drawbacks to using tapes. Data stored on tape is inherently sequential, so random or directly processing individual records is impractical. Tape volumes are not mounted until they are required, as opposed to most DASD devices, which tend to remain more or less permanently mounted. This implies human intervention, which causes both a time delay in mounting the tape and a greater potential for error in handling than is typically encountered with DASD. A tape can only be accessed by the job that called for it, unlike a file on DASD which can be shared by multiple processors at the same time.

There was a great need for a mass storage system with a large data storage capacity which would be under system control, and have the data organization flexibility of DASD. It needed to have "current" DASD transfer rates and a cost per megabyte of storage closer to that of tape than DASD. When IBM announced the IBM 3850 Mass Storage Subsystem, it met these requirements with sophisticated technology extending the concept of virtual storage to the I/O components and providing the capacity of a tape library. Availability and mounting of the volumes was under the control of the operating system with the same variety of methods of data organization available on DASD. Even multi-volume data sets could be used. [Ref. 12, p. 41]

The MSS records data on 2.7" wide by 55" long magnetic tape contained within cylindrical cartridges, 3.5" long with a 1.9" diameter. Two of these cartridges are referenced as one 3330V (V for virtual) volume and hold 100 megabytes of data, in the image of a 1974-vintage 3330 Model 1 disk volume. The MSS consists of a Mass Storage Facility with Data Recording Devices, Data Recording Controls, and Mass Storage Controls with Accessors which take the cartridges from the honeycombed storage walls to the Data Recording Devices, returning them when finished. There are several possible configurations of the system. These vary from a minimum capacity of 35 gigabytes of

data (equal to 350 3330-1 volumes) to 472 gigabytes (equal to 4,720 3330-1 volumes) in the maximum configuration. [Ref. 13] The NPS system was a model A02 with four Data Recording Devices, two Data Recording Controls, one Mass Storage Control, with a capacity of 2,044 data cartridges which equates to 1,017 virtual volumes with a capacity of 101.7 gigabytes. (Ten cartridge cells were reserved for operational considerations and maintenance.)

On the MSS, data is *staged* from the IBM 3851, onto real IBM 3330 or 3350 disk storage in eight cylinder segments for as long as it is needed. (See Figure 2) Then, the data is *de-staged*, and the physical DASD can be used for eight more cylinders of user data. When the data is staged, it can be shared by more than one MVS job, as can any regular DASD data set. MSS can also be used for the VM user mini-disks. In June 1987, the Center's MSS had 75 volumes for online, time-sharing users of VM/CMS and 314 volumes for batch processing under MVS SP.

Mass storage volumes are defined in groups with a name and an owner. After the group is defined, more volumes may be defined to the group as desired. The Center's groups were primarily by department, with some departments having multiple groups. With the MSS-provided ability to assign mass storage volumes to groups, the storage manager had some control over the use of the volumes. Since the inventory data set group record contained the information for the group, allocation parameters could be specified for blocksizes and space allocations for data sets. Individuals did not have to specify their DASD requirements. The default parameters were used, whether or not they were optimum. [Ref. 12, p. 5] IBM recommended using naming conventions for improving control of application data sets and as future guidance to application programmers. [Ref. 12, p. 41] NPS users generally used the defined naming convention but did not always follow the recommendation to catalog all data sets. If the user did not follow the naming convention, his data set could not be cataloged. Cataloging finally became accepted by all users one year after they were told that it was required.

The concept of volume ownership by a user group dates back to the days of removable DASD. As a physical security measure, the volume could be removed from the Computer Center and stored elsewhere. To maintain reliability with today's technology, at such high access rates, DASD cannot be removable. With the capacity of DASD in the 1980's, it is extremely costly for a particular group to own its own volume. The largest IBM removable volume, IBM 3330-11, which the MSS simulates, holds 200 megabytes. A double density IBM 3380 Model E holds 2.5 gigabytes, equivalent to

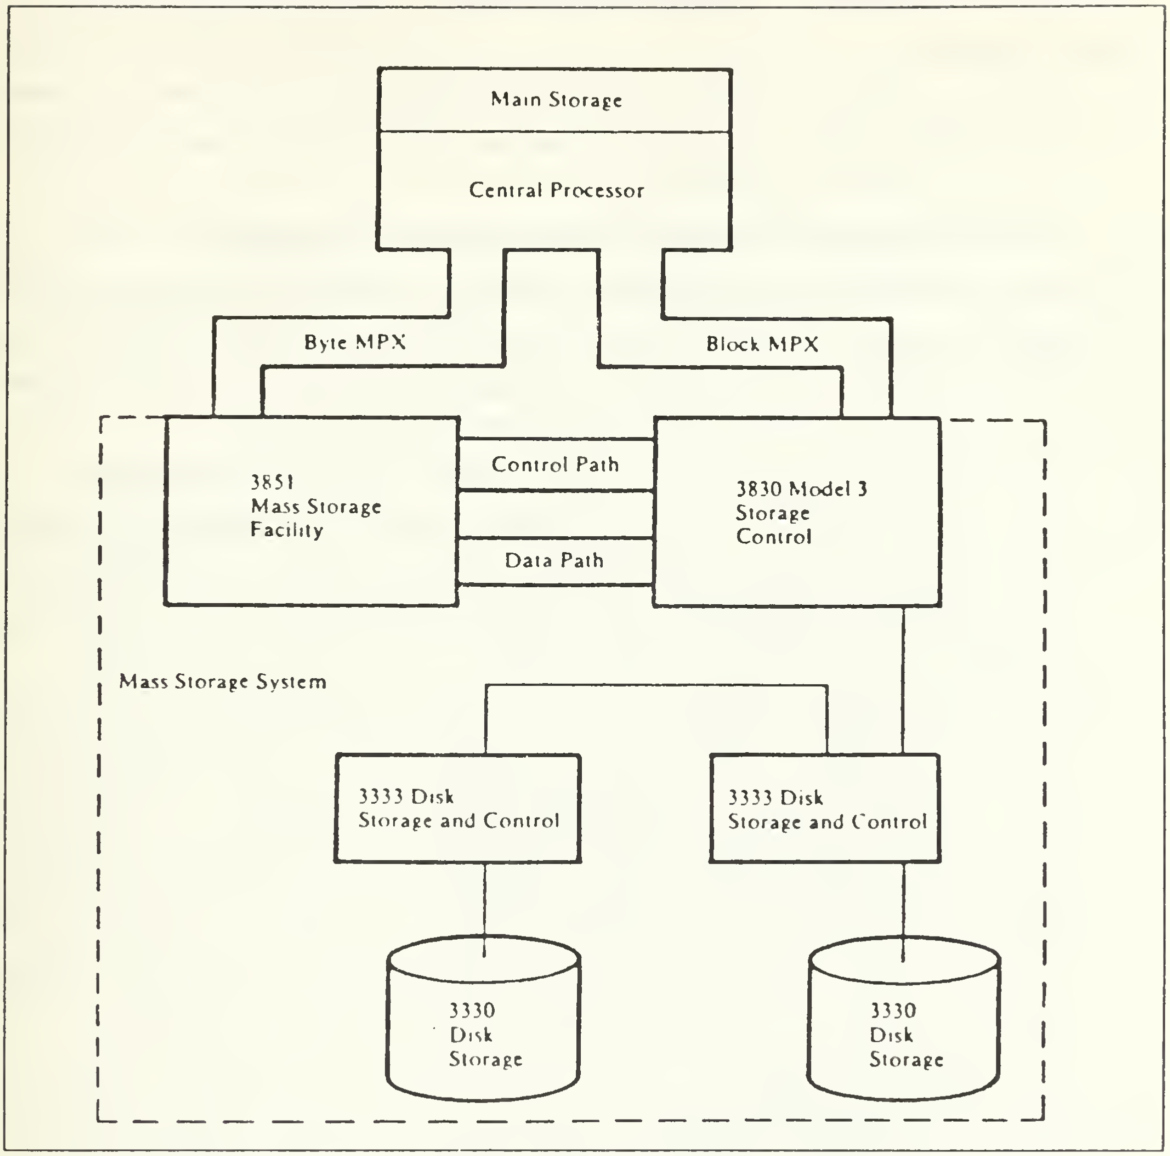

Figure 2. Mass Storage System Configuration

twelve IBM 3330-11's. New methods of management and control must be established in this environment. [Ref. 6]

With an average access time, after staging, comparable to IBM 3350 disks, and a data transfer rate of 806,000 bytes per second, the MSS cannot begin to keep up with modern processors. [Ref. 14] The speed of processors has increased significantly in thirteen years, but the Mass Storage Controller (MSC) processing speed of 30 to 45 orders per minute has remained fairly constant. Currently one CPU can send orders to the

MSC faster than the MSC is able to process them. "When most installations were running processors of the 370/145 to 370/168 class, the MSC could handle requests and commands from multiple CPUs." [Ref. 11] The 308X and 309X classes of processors easily exceed the MSC order rates. Processors running jobs that require MSS data will be severely limited by the 3850 speed.

DASD and control unit speeds have also increased significantly. DASD transfer rates over 3 megabytes per second have become a requirement to keep today's CPU running efficiently. As with the MSC, the 3830-3 Staging Adapter (SA) and the staging DASD continue to transfer data at the nominal rate of about 800 kilobytes per second. As demands increase, the effective data rate of the SA's is reduced.

Many computing facilities installed the MSS as a low-cost storage device. At the time, it was a good choice of storage media for large quantities of infrequently-used data. Although large and inexpensive, this seemingly infinite, virtual DASD had hidden costs. To keep the 3850 Subsystem running properly, the facility required trained systems programmers to install, maintain, and support it. Many installations required a person to spend full-time on the MSS--learning the product; managing data spaces; and recovering from component, volume, and subsystem failures. This expertise came from working on the Subsystem, and from IBM classes and workshops. As MSS reliability improved with resultant fewer outages, the recovery skills of systems programmers were exercised less frequently. When problems occurred, this made MSS recovery a longer and generally more difficult process. Users compounded this problem by storing many production data sets on the MSS. Occasionally, when the MSS would not be available and users needed this data, they would have to wait for the systems programmer to complete the problem determination and recovery. During these times, little productive work was done in the installation, and the MSS outage was extremely visible. This predicament could be avoided by migrating MSS data to DASD and tape, and keeping the MSS out of the critical path of the installation's production jobs. [Ref. 11]

Besides the emergency need for the systems programmer to identify problems and recover from them, much time was needed on a continuing basis. The Mass Storage inventory and journal data sets required backup and attention from the systems programmer. A duplicate set of the MSS tables must always be available in case of failure of the primary tables. [Ref. 12, p. 5] Switching tables and recovery from table failure is not a trivial matter as was learned more than once. In April 1987, the Center experienced a recurring failure and a system outage of approximately eight hours on one day.

Where the MSS was the best solution in 1976, with faster processors and other I/O devices, the MSS will not be able to satisfy users who need a large amount of storage at a reasonable cost, online to multiple fast processors in the 1990's. With IBM's announcement of withdrawal of support of MSS on future systems, MSS users had to begin migrating Mass Storage Subsystem data to other storage devices.

The report, **3850 Mass Storage Subsystem Migration Planning (GG66-0208-0)**, published in August 1985 [Ref. 11] described several migration strategies which could be used to replace an MSS in an orderly manner. This document does not discuss evaluation of whether to replace the MSS or not, but offers several approaches for the task. In 1985, some installations were quite content with their use of the MSS. If they had developed the needed recovery skills and were pleased with their use of the MSS, they might see no reason to change the way they store and use the data in the MSS. If it was satisfactory for their application, they wanted to know why they should migrate to something new. Until the 3090 announcement, this attitude was understandable.

In February 1985, part of the IBM announcement package for the 3090 processor was a letter stating that IBM did not intend to support the attachment of an MSS to any IBM processor beyond the 3090 Processor Complex. [Ref. 1] Even for the satisfied user, IBM recommended consideration of alternatives to MSS. The results of this review should be a plan to replace the MSS with DASD and tape that would connect to any family of processors. NPS must be prepared for further developments and make transitions and new purchases in an incremental fashion to lower costs and the impact of radical changes to the users of the NPS Computer Center.

Over the last ten years the cost, floor space, and data density of DASD have made this type of storage more competitive with MSS. A combination of 3380E DASD and 3480 tapes can provide an excellent alternative to the MSS Subsystem. They provide the solution to both the recovery skills problem and the MSC transfer rate problem. [Ref. 11] They represent current technology and allow for growth to future developments.

IBM recommended that the first step to knowing what to do about the MSS is an analysis of its use. Classifying the data would tell you what should be done. "You can't decide where to go without knowing where you are." [Ref. 11] The three categories suggested by IBM are active, inactive, and a combination of the two. Inactive is defined as data that is either system managed or a copy of user data, created by a storage management product such as IBM's Data Facility Hierarchical Storage Manager (DFHSM) or Data Facility Data Set Services (DFDSS). This data would not be directly accessed by an end user, if at all. In 1985, a number of users with DFHSM installed used the

MSS for Migration Level 2 (ML2) storage. IBM defined active data as data that is accessed directly by the end user, either in a test or production environment. Data in this category would be referenced quite frequently, with, perhaps, some production job dependencies. Files belonging to staff and faculty members that have not been used in a long time (maybe years) are inactive, although the owners would want the files easily addressable by an MVS job stream. IBM's recommendation [Ref. 11] was to migrate all active data to something other than the MSS. At that point, how long the MSS was used in the Center would be up to the customer. This migration would take time and there could be a few interim steps and hardware configurations planned before the final data storage configuration is achieved. The options recommended were to move all data, move only the active data and wait until the rest of the data was obsolete, or if it contained only inactive data, wait until it was obsolete, then remove the MSS. Each option included changes in hardware and software. The configurations recommended by IBM included 3480 tape and 3380 disk hardware, and DFHSM as a storage manager.

In most 3850 installations, a large percentage of the workload depends on the MSS being continuously available. Any MSS outage caused missed deadlines and delays to many users. For the NPS, an outage affected nearly every user, with the individual accounts (VM mini-disks) on MSS and the MSS an integral component to MVS. The Reliability-Availability-Serviceability (RAS) design and significant performance advantage of 3380 DASD are superior to that of MSS and its components. Under MSS, with the stage-destage rate of approximately 125 kilobytes per second, a user must wait at least two seconds after the virtual volume has been mounted for a data set to be opened. If the data set is 16 cylinders or more, the user must wait sixteen seconds. A maximum of four processors can connect to an MSS, but any one processor can attach to only one MSS. In addition, only three processors can be connected to any Staging Adaptor, or Staging Adaptor pair using a redundant path. [Ref. 11] With 3380 DASD, eight processors can be connected to any device plus an alternate path. In a 3480 tape subsystem, the A22 control unit pair can have four CPUs connected, each with an alternate path. Although connectivity was not a problem for NPS, it would have a significant bearing in a large MSS installation, such as an insurance company with many years of data on the MSS. Upgrading to a new version of the operating system or to a new processor family would be inhibited by the MSS.

In IBM's **3850 Mass Storage Subsystem Migration Planning**, [Ref. 11] five configurations are proposed as a replacement for the MSS. All of the configurations include 3380E DASD, and some include 3480 tape drives. Four require some form of storage

manager, such as DFHSM. Although not a necessity, IBM recommends installing it, as it would provide an automated storage manager to replace much manual effort. The benefits derived from its use are both immediate and long-term and are discussed in Chapter IV.

The first three configurations recommended by IBM are

1. All active data moved to 3380 DASD, inactive data moved to a combination of 3380 DASD and 3480 tape. A storage hierarchy of 3380 DASD and 3480 tape with all active data moved to DASD only, DASD for DFHSM Migration Level 1 (ML1) storage, and tape for DFHSM Migration Level 2 (ML2) and backup storage.

2. Active data and inactive data moved to a combination of 3380 DASD and 3480 tape. A variation of the first configuration, but assumes that the movement of some of the active data to 3380 DASD is not justified, due to the size or frequency of access. A storage hierarchy of 3380 DASD and 3480 tape with all active small and intermediate files moved to DASD, large files moved to tape, DASD for DFHSM ML1 storage, and tape for DFHSM ML2 and backup storage.

3. No storage management product implemented, and the data currently residing in the MSS (both active and inactive) moved to a combination of 3380 DASD and 3480 tape. Management of all devices done manually. In some cases, it would be necessary to replace storage management functions currently being performed by MSS utilities such as SCRDSET, or MSSE functions such as System Initiated Scratch. There are no equivalent IBM alternatives other than DFHSM. A storage hierarchy of 3380 DASD and 3480 tape with all active small and intermediate files moved to DASD, large files moved to tape, and with no storage management product installed. [Ref. 11]

The other two configurations contained MSS as an interim configuration, only. Basically they are the same as those above, but the inactive data is left on MSS until it is obsolete. The objective of the migration is removal of the MSS.

In order to estimate the new DASD requirements, a study was done of the MSS volume assignments and utilization. Space analysis was performed to determine what percentage of the MSS volumes were actually used in order to determine the amount of new hardware needed. The results are included as Appendix A. The use of 3480 tape drives and DFHSM for compaction reduces the floorspace necessary for more 3380 DASD.

The new releases of DFDSS and DFHSM which were provided on the Custom-Built Installation Package (CBIPO) provide the means to move the data easily and manage it effectively in the new environment. In order to use these products effectively, the data sets to be migrated had to be cataloged in Integrated Catalog Facility (ICF) catalogs. With this requirement as a prerequisite for subsequent steps, the beginning of the migration at NPS was the catalog conversion since the primary user catalog was

of the old style control volume (CVOL) catalog. Chapter V will describe the specific steps in the migration process.

IV. STORAGE MANAGEMENT NEEDS–HOW DFHSM SATISFIES THEM

In order to understand what is needed in a storage manager for DASD, one must first understand the tasks to be done. Storage management tasks fall into three categories: device, space, and data management. Device installation and maintenance are covered sufficiently by Device Support Facilities, ICKDSF, an IBM utility, used by Computer Center's systems programmers. Management tasks are primarily space and data set management.

A. SPACE

Space is allocated for data sets and released when data sets are deleted. This requires some control to ensure data sets are allocated to proper volumes and deleted when no longer needed. Sufficient space for allocating new data sets and extension of existing ones is critical. To ensure this, space must be used efficiently and effectively. If unblocked, 80 byte records use less than 15% of a 3380 track. Full track blocking provides for maximum use of DASD capacity. Since there is no way at present to have default block sizes, the user must specify them. Optimum block sizes depend upon the track capacity of the specific device. If data sets are to be moved from one device to another, standards should be developed which effectively utilize the capacity for the types of devices used. Although, a block size of 19069 uses 100% of an IBM 3350 track, it uses only 80% of an IBM 3380 track. For data sets which are stored only on an IBM 3380 (or 3380-image) device, a block size of 23,476 (half track) is optimal, allowing space utilization of 98.9%. [Ref. 15 p.146] For data sets to be moved between IBM 3350 and IBM 3380 units, a block size of 9076 uses 95% of both units. [Ref. 6]

Allocation of space by cylinder was the default for user data sets on the MSS. However, this wastes space on an IBM 3380. The documentation given to the users for the migration of data sets from the MSS to IBM 3380s stressed allocating in blocks, not cylinders. This was something new so it required some learning on the part of most users. Allocation in blocks assures good capacity utilization regardless of device type. Cataloging and eliminating the use of job or step catalogs reduces the chance of having duplicate data sets in the system.

For the actual space on the volumes, Data Facility Hierarchical Storage Manager (DFHSM) is irreplaceable as an aid in the areas of relocation (migration and recall),

conversion, movement, retirement or archiving, and deletion. This is the product recommended by IBM for the migration from MSS. [Ref. 11] It controls the amount of data on a volume, deletes obsolete data by age, creates backup copies, compresses data sets and moves them off the primary volumes, making room for more currently used data sets. When a volume becomes highly fragmented, the storage administrator can use DFDSS, an IBM utility, to rearrange the data sets to make available larger contiguous spaces for re-allocation. Reorganization of the free space on a volume does not make more space. This brings us to the discussion of the data sets themselves.

To manage the available space, there must be system-wide procedures for migration of data sets from the high availability DASD. DFHSM supports a hierarchy of levels of access. Primary volumes contain data for current use. The Center allocated ten 3380 volumes as primary volumes (12.6 gigabytes). In order to have space for new allocations, space management runs daily with parameters specified to migrate any data set from a primary volume over 60% full to a Migration Level 1 (ML1) volume in a compacted form if it has not been used in ten days. The Center allocated three 3380 volumes for ML1 (3.78 gigabytes). From ML1 volumes, data sets not used in 25 days will be migrated to Migration 2 (ML2) volumes. Three times, the ML1 volumes have filled up completely causing migration to ML2 volumes. Until the ML1 volumes fill to 80%, all of the data sets, no matter how old, are available to the users with no operator intervention. This has given availability of data on ML1 volumes for three months or longer. Two hundred 3480 tapes were acquired to be ML2 volumes. After 14 months, 76 ML2 tapes are between 50% and 99% full. Removing unused data sets, which migration does, is automated under DFHSM. Without DFHSM, the storage administrator would need to do this task.

B. DATA MANAGEMENT

The many areas of data management range from the creation to the deletion of data: creation and classification; control as related to identification and location, access (authorization, availability, performance), monitoring usage, standards; relocation as in migration and recall, conversion, and movement; retirement or archiving; and deletion. Naming conventions and aliases aid in control and classification, identification and location.

Who is responsible for what tasks in the management of data sets? Ideally, the application users should be responsible for logical data, the system should be responsible for physical storage, with the storage administration group serving as a policy/control

interface between them. [Ref. 6] In an ideal environment, the application users should only have to be concerned about the logical view of their files. This means that tasks such as backup, recovery, space availability, and volume clean-up are the responsibility of someone else. In the environment of personally-owned volumes, the user was responsible for these tasks. For the system to be responsible for physical storage, there must be some interface between the system and the users. An IBM speaker at SHARE called this the Storage Administration group. "Studies conducted in 1982 and 1983 indicate that it took one person for every ten gigabytes of DASD to perform the storage management tasks. At the average compound growth rate of 45%, even the smaller installations, ... will require a large number of people just to perform the DASD management tasks." [Ref. 6] In addition to being the interface between the application users and the system, this group would be responsible for all areas of DASD management such as policy definition and control; device selection, installation, and usage, space allocation, capacity utilization, capacity planning, service level management, installation standards, performance, availability. Additional tools and techniques need to be developed and used that allow a storage administrator to manage large quantities of DASD (100-300 gigabytes). With a compound growth rate of 30-45 percent, new hardware is always part of the solution to storage management problems, but tools which automate as many of the storage management tasks as possible are needed to minimize the personnel requirement of storage administration. Standards, especially data set naming standards, affect the ability of storage administration to automate the management tasks via software and standard procedures.

The need for additional DASD capacity is always present. However, there is a point at which additional DASD capacity may not solve the storage management problem. For further information regarding the balanced system concept, the reader could refer to **Capacity Planning, Basic Hand Analysis** by L. Bronner, IBM publication number GG22-9344, or **Balanced Systems and Capacity Planning** by R. J. Wicks, IBM publication number GG22-9299-01.

C. SUMMARY

In Chapter II, it was stated that a mass storage system should provide:

- Relatively fast access
- Data access compatible with systems software and access methods
- System accessible storage media
- Technologies which can be enhanced to provide long-term storage solutions

- Cost effectiveness not only based on price but on operational and environmental measures. [Ref. 2]

With DFHSM, access for a data set used within the last seven days is at the rate of 3 megabytes per second, as fast as access to any data set on a 3380. If the data set has been migrated to ML1 and, therefore, compacted, it must be decompacted when being moved from one 3380 to another. Although we don't know the rate of decompaction, the amount of time is negligible. If the data set hasn't been used for a long time and has been migrated to ML2, it takes longer to retrieve, since operator intervention is required to mount the correct IBM 3480 tape. From that point, the system operates at the speed of the high-speed tape and the high-speed DASD, viz., 3 megabytes per second. After recall from either ML1 or ML2, it will remain on the 3380 as long as the user accesses the data set at least once a week.

DFHSM uses standard IBM utilities to do the functions it requires. When these utilities are improved, the improvements will be automatically a part of DFHSM. Along with using the latest in software, the hardware which can be used is IBM's latest. DFHSM has been announced as a strategic component of IBM's MVS/ESA (Enterprise System Architecture). That makes it clear that support for new hardware and software will be a part of DFHSM. Currently it supports the following devices: 3330 direct access storage devices, models 1 and 11; 3350 direct access storage devices; 3375 direct access storage devices; 3380 direct access storage devices, models A04, B04, AA4, AD4, BD4, AE4, BE4 and all the J and K models; 3850 mass storage system; 3420 magnetic tape units; 3430 magnetic tape units; and 3480 magnetic tape subsystem. With this support and the MVS/ESA announcement with DFHSM as a strategic component, the questions of compatible media and long term support are satisfied.

V. IMPLEMENTATION

A. CATALOG CONVERSION

Implementing the DFHSM program required that all of the catalogs be Integrated Catalog Facility (ICF) VSAM catalogs. The catalogs were old style VSAM, with the user catalog the older style control volume (CVOL) form. The first step in the DASD migration process required converting all the MVS catalogs to ICF VSAM catalogs. The system master catalog conversion was first and it went smoothly. Each user catalog has an entry in the master catalog. Along with references to the user catalogs, the master catalog contains "aliases" which refer to the high level index, or first segment, of the data set name. An alias points to the user catalog where a data set with that high level index will be cataloged. The use of aliases helps enforce the use of some naming conventions. [Ref. 16] The master catalog is password protected while the user catalogs are not. If the user follows the established naming convention, the data set can be cataloged in the proper catalog. Otherwise, the system will try to catalog it in the master catalog. The operator cannot give the password, therefore the data set will not be cataloged if the established naming conventions have not been followed. When the data set is not cataloged, DFHSM indicates this fact on its daily space management report. Eleven days after creation, the data set will be deleted. The naming convention in effect at NPS is a simple one. The first segment correlates a defined alias to the catalog in which the data set is to be cataloged. It must be correct in order for the data set to be cataloged. For the basic user, the second segment contains a code of an alpha character indicating the user category (student, faculty, computer center staff, and others), followed by the user number, thereby identifying the owner of the file. Some other special users have their own first level index to put files in separate catalogs. For them, the second segment has specific identifiers to establish ownership of the data set.

The IBM conversion utility failed on the second catalog, the catalog for a strategic School function. After restoring the VSAM files for the third time from the backup, normal recovery procedures were considered. There was no ongoing procedure for a frequent, regular backup of these strategic VSAM files. This vulnerability was corrected by instituting a weekly backup of these files in order to recover from future failures. IBM's support was requested on the conversion of the second catalog. They sent documentation of a new function of the IDCAMS utility, DIAGNOSE, which analyzes a

VSAM catalog for errors and should be used prior to conversion. IBM felt that errors in the catalog had caused the conversion utility to fail. This hypothesis could not be confirmed because our restoration of the catalog could have caused the incongruencies that existed at that point. This DIAGNOSE function was run against the remaining VSAM catalogs prior to conversion. None had errors. Since the conversion utility would not work on the second catalog, each data set had to be EXPORTed (copied) to a sequential file for backup, deleted, allocated in the new catalog, then IMPORTed (copied) from the sequential file. This was a much more time-consuming process than experienced when using the IBM utility provided for the conversion process.

The catalog in which all the user data sets were recorded was of the older style CVOL catalog. It was converted to an ICF VSAM catalog during the last week of December 1987 with no problems. The major difference to the users, between the old style catalog and the new one, was the deletion of an outdated utility (IEHPROGM) which did not reference the catalog. The users were notified of this fact several weeks prior to the conversion of this catalog. The old utility was no longer required. Another utility (IDCAMS) did the same functions but the users did not want to change JCL that worked...or JCL *they thought* worked. (After the conversion, 1100 entries were removed from this catalog for data sets supposedly deleted by using the old utility. However, the old utility did not delete the entry from the catalog. It merely scratched the data set.) This was the beginning of the visible reluctance of the users to accept the changes that were to come. Along with the requirement to use the IDCAMS utility instead of the IEHPROGM utility, the Computer Center published an article containing detailed instructions and JCL for use when cataloging data sets. After the data set was created, the user need only refer to it later with the DATASETNAME (DSN) and DISPOSITION (DISP) parameters. Using the data set entry in the catalog was the first step for an easier transition for later changes. Many users ignored this recommendation.

According to IBM, [Ref. 11] probably the most difficult part of the MSS migration will be the JCL modification necessary to direct all new allocations away from the MSS. There are no IBM products available to perform the changes. Except for procedures such as the FORTRAN compiler procedures, etc., it was recommended that all JCL just reference the cataloged data sets with only DSN and DISP parameters on the data definition (DD) statements. For allocation, the generic UNIT=SYSDA replaces UNIT=3380 and a specific VOL=SER=nnnnnn from the pool of volumes.

Installing the hardware, the additional 3380s and controllers and the 3480 tape drives, was the next step. The hardware installation was originally scheduled for the

spring break, a long weekend between quarters at the end of March. The lack of courses with accompanying assignments for the students during this short period would have allowed the Computer Center the luxury of having the computer completely down for a few days. Unfortunately, this schedule could not be met because of procurement delays.

If the hardware had been installed at the end of March, the Custom-Built Installation Package Offering (CBIPO) for MVS with DFHSM could have been installed four or five weeks earlier. DFHSM then could have been used *gradually* on groups of files, beginning with the ones belonging to the Computer Center staff. This scenario would have allowed the writing of the "cookbook" technical newsletters and procedure files to aid the users in handling their data sets. If the Computer Center staff had been able to use the product for a short while, most of the problems which we experienced might not have occurred. A gradual changeover would have made for a smoother migration. Though some users resisted, the changes would have been more transparent had the entire Center's staff been more able to help. The primary differences concerned allocating new files with block instead of cylinder allocation and using the IDCAMS utilities. It seemed that some of the Computer Center staff objected to the change as well as the less experienced users.

B. USER DATASET EVALUATION AND BACKUP

Some users wanted to backup their data from the Mass Storage Subsystem (MSS) prior to the migration. Defense Manpower Data Center (DMDC) moved some of their data from the MSS to DASD owned by them. Generally, DMDC data residing on the MSS was of a backup, archival nature. Therefore, most data was moved to tape. The initial attempts at producing these backups were not very successful. DMDC's Full-Virtual-Volume-to-tape jobs (using DFDSS, an IBM utility for dump, restore, or copy operations) required six to eight hours for one MSS volume. Investigation revealed that MSS was staging the data sets eight cylinders at a time! To overcome this problem, Mr. D. Norman, Manager of the Systems Support Group for the Center, wrote a small assembler language program which accessed the MSS Communicator at the SVC level, mounted and staged a complete volume, then invoked the DFDSS copy program. The program then released the MSS volume. Done this way, the backup process was reduced to 10 to 30 minutes per volume.

The identification of obsolete, deletable data was left to the user. Many phone calls were made to the owner of each MSVGP group. The owner, alone, knew the value of the data sets on the virtual volumes. It was assumed that each owner had done the

requested evaluation by the specified deadline. If the data sets were not cataloged, as required, it was assumed that the user did not want the data set. Every cataloged data set on each MSS volume that the owner did not personally delete or ask us to delete was migrated. The users, if available, were contacted for confirmation that they did not want each uncataloged data set. Very few of these were subsequently cataloged and moved.

According to IBM, "the migration from the MSS is not going to be without some cost. The Space Manager is going to have to do the majority of the work, but the co-operation of the end-users will be important to the success of the migration because they will have to do some of the work. Even if the majority of the JCL changes, data deletion, and data movement can be done by a Production Control group, a Storage Manager, or a combination of the two, there will be a required participation by the end-user community. JCL that exists in individual data sets must be changed, obsolete data sets must be identified, and some of the data movement will have to be done by the end-users. It will be useful to gain the support of the end-user community early in the planning cycle so that they are aware of the work that must be done." [Ref. 11]

Upon request, Center users were assisted in creating their own backup copies of their data sets prior to the migration. Earlier backups from 3850 would not be able to be restored once that hardware was removed. Some users fought the changes, although the changes were few in number. Generally, JCL was simplified in the process.

C. SOFTWARE INSTALLATION

The Custom-Built Installation Package Offering (CBIPO) for MVS was installed during May 1987 after the installation of the 3380s and 3480s. This task required about four weeks of full-time effort by a senior systems programmer. This CBIPO contained the base MVS operating system with no major systems changes, except the addition of DFHSM. New releases of several utilities were included which were needed to support DFHSM and upgrade the system components to current levels. A CBIPO with more new functions and changes would have taken even longer to install.

D. MIGRATION

The VM mini-disks were moved from MSS to 3380's over Memorial Day, May 1987. The VM Systems Programmer did all of the copying, making the change virtually transparent to the users.

The procedures for implementation of DFHSM were set up and the migration begun. Volumes of the Mass Storage Subsystem (MSS) were migrated to DFHSM Migration Level 2 (ML2) volumes on 3480 tapes. The initial migration was begun on a

Sunday, June 21. Recall was tested on various types of data sets from different ML2 tapes. Each task worked perfectly. As we attempted to recall data sets on the loaded system during the following workdays, we found that the tape drives were never allocated to DFHSM for the recall of the needed data sets. After much research and many phone calls, IBM responded with an "undocumented" parameter which makes it possible to define the DFHSM 3480 tape drives as a different unit from the other 3480 tapes drives. The migration continued, with the majority of the data sets being moved over the Fourth of July weekend.

In recalling migrated data sets, we ran into two categories of data sets that caused problems: (1) direct access data sets (created by FORTRAN programs) and (2) data sets which could not be reblocked upon recall. The direct access data sets created by SAS programs had been moved prior to the migration. (SAS is a Statistical Analysis System from the SAS Institute, Cary, North Carolina.) A Center staffer consulted with another university which had been using DFHSM for several years and inquired about the effects upon SAS data sets. The reply was that DFHSM handles the data sets created by SAS as long as they are migrated from, and recalled to, the same type of unit. This would not be true during our migration, from the 3850 with its virtual 3330 volumes to the real 3380 volumes. A known procedure was used to move SAS data sets to 3380s prior to the actual migration. There were about 1,500 other direct access data sets on the system. The IBM documentation states that DFHSM will handle direct access data sets properly. [Ref. 17] After discussions with IBM personnel, we assumed that our direct migration would work. In fact these direct access data sets were handled in the same way the SAS data sets were, that is, fine when migrated from, and recalled to, the same type unit. The migration proceeded. Afterwards, these direct access data sets all had to be recalled first to a 3350 volume simulating a 3330-1, then moved with DFDSS, a utility for copying data sets, to a 3380 disk prior to use and control by DFHSM. IBM assured us that future documentation would be clearer on this point.

The second problem, data sets which could not be reblocked upon recall, was solved by IBM. A parameter on the DFHSM control procedure was changed for one CPU from CONVERSION(REBLOCKTOANY) to NOCONVERSION. If the job calling for the data set failed when running on SY2 (the 3033U CPU), the user could contact a Computer Center staff member with TSO access. TSO runs on SY3 (the 4381 CPU) where DFHSM was set up with the NOCONVERSION parameter. Or, the user could resubmit the job, specifying that it should run on SY3. This problem only occurred on the first recall of the data set after migration from the MSS.

From the first of July 1987 until the first of September 1987, no other problems were encountered. In early September, two unusual partitioned data sets (unformatted) could not be recalled from the DFHSM ML2 tape. Other data sets created the same way had been recalled. The IBM Support Center defined the problem as a missing end-of-file-marker on the members (or a member) in each of these two partitioned data sets prevented recall. IBM accepted this problem from us and a few other users to make an immediate code change for DFHSM. It will no longer allow data sets such as these to be migrated or backed up in the first place, and will give an error indicating that these data sets are not standard and not covered by DFHSM.[1] Luckily, the owner of the two data sets, a doctoral candidate, was able to recreate them.

In mid-September, a DFHSM ML2 tape would not allow the recall of any of the data sets on it. The label had been damaged by an operator re-labeling it after DFHSM had written files on the tape. A program was written to copy the DFHSM tape to a new tape, a record at a time. Only the first file was lost from the tape. The procedure was documented and the program was saved for similar future recovery situations.

E. ONE YEAR LATER

After monitoring DFHSM for a year, observing the storage management functioning well, the Center is quite pleased with DFHSM and all the storage control it affords. There may be some tuning which still needs to be done, but not as much as was imagined when the project was begun. The original estimates of how much space should be allocated to each level appears to be still appropriate. On August 22, 1988, one primary volume was removed from DFHSM's control in order to use it for another purpose. This caused the other volumes to fill. One migration parameter, how long data sets would be allowed to stay on the primary volume since the last reference, was lowered from ten to six days. After approximately two weeks, with much forced migration by the author, it appears that DFHSM has evened out the load.

At first glance, the migration would seem to be a step backwards: from an automated mass storage system entirely under system control to one with operator intervention required at one level. However, the annual cost savings are substantial and considerably more space is available for the users.

Managers of automated data processing installations have to stay in touch with trends in technology. The changes that will come with IBM's Enterprise System

---

1 APAR Number OY08276, December 4, 1987

Architecture (ESA) will be dramatic. The user's view of storage is changing from the physical to the logical. After a data set is created and cataloged, its physical location is not important to the user. The system can move the data set around to maximize the use of the physical DASD. DFHSM does this. IBM has stated that DFHSM will be a strategic product in ESA. The Computer Center is now well-positioned for the future developments.

A snapshot profile of storage utilization was taken on September 2, 1988, 8:41 a.m. At that time, 119 users had data sets on primary and Migration Level 1 storage. Table 2 shows how much space (in tracks) and what percentage of the total that the top ten percent of the users had allocated. It comes as no surprise that they used 76.5% of the allocated primary storage and 19.3% of the space used on ML1 storage. The top user, a student in the Air-Ocean Curriculum approaching graduation in December, had 49.5% of the allocated primary space. This is equal to nearly 20 old MSS volumes plus nearly two more for the data sets on ML1 storage. He also had 155 more data sets on Migration Level 2 (ML2) on 3480 tape. He is using the files on primary storage on a continuing basis or they would be migrated. He and the second largest user, a doctoral candidate, each have 137 files allocated on primary and ML1 storage. The second user has 187 more data sets on ML2 storage. Prior to the migration from the MSS to DFHSM, a user was limited in the amount of space and the number of data sets held on public storage. At this time, the Center has not defined limits to the amount of space a user can allocate. This policy will be reviewed periodically.

Since user S2310 had so much space allocated, the author looked at the utilization levels. On some of his data sets, he was using only one-third of the allocated space. These are large data sets, therefore the amount wasted was considerable. Besides keeping an eye out for his data sets which could have excess space freed manually, the author counseled the user on his job control language. He was only too happy to add the RLSE parameter to his space request to release any space not needed by the jobs that he was running. He was using job control language given him by another user and really had only a rudimentary knowledge of what the job control language was doing and no idea what other options were available to him.

The snapshot contains information which would be beneficial to the Computer Center as an aid in monitoring users. Because of this, the author wrote a program to obtain this information. It can be run as often as is necessary, but current plans are to run it monthly. The first run, approximately three weeks after the snapshot view, showed that user S2310's usage had dropped from 49.5% of primary storage to 38.6%.

He is still using a large amount of storage, but there is 100% utilization of his data sets now. The first run of the program showed that the top 10% of the users were using 77.19% of primary storage. Six of the users from the snapshot were still in the top 10% three weeks later. Two users on the snapshot (CC08 and C0037) belong to the Computer Center's Accounting Office. The author's program combines them into one user. Because the files used by the Accounting Office are used continuously, they will not be subject to migration, but will be moved to another volume not under control of DFHSM, when such space becomes available. The program also combines the information about other users with more than one user ID in order to have a more definitive description of system usage by user.

Table 2. TOP 12 USERS OF STORAGE

USER NUMBER	PRIMARY		ML1		TOTAL
	Tracks	Percent	Tracks	Percent	
S2310	83,631	49.5	7,642	4.5	91,273
N0196	9,566	5.5	11,903	6.9	21,347
CC08	5,440	3.2	51	.0	5,491
C0037	5,170	3.0	28	.0	5,198
S3056	5,111	2.9	5	.0	5,116
N3945	5,093	2.9	1,030	.6	6,123
F3862	3,538	2.1	1,687	1.0	5,225
F3964	3,400	2.0	487	.3	3,887
F5008	3,149	1.8	36	.0	3,185
N3958	3,065	1.8	853	.5	3,918
F3910	2,497	1.5	5,324	3.1	7,821
F3902	531	.3	4,014	2.3	4,545
Total	130,191	76.5	33,060	19.2	163,129
Users are in descending order by space allocated on primary volumes.					
Percentages are of the total allocated space for that category.					

The Center will continue to monitor the utilization of space to see if the correct amount is allocated to primary and ML1 volumes and if two hundred 3480 tapes will handle the data for two years as in the original forecast. The parameter specifying how long data sets remain on primary storage has already been changed, but only after 14 months.

The challenge will be to decide whether to limit the users and, if so, how. The above table shows that the users at the bottom of the table, with amounts of data on ML1 comparable with other users' primary storage, have not been working with those data sets for over a week. This indicates that DFHSM is doing what was intended.

APPENDIX A. SPACE AND USAGE ANALYSES OF MSS VOLUMES

The following data on the total space in user data sets were acquired at approximately the same time of year (the second week in May); in 1986 and 1987, by a single run each with an IBM utility on the MSS and, in 1988, by listing all the data sets on the relevant 3380 disk and 3480 tape volumes. The data concern user data sets existing on MVS volumes on May 5, 1988 with MVS CPU usage during the academic quarter, January 4 through March 24, 1988, as reported by Duquesne's Billing Database Facility (BDBF) accounting package.

A. SPACE

Summary statistics are shown in Table 3.

Table 3. COMPARISON OF THE USE OF MASS STORAGE—1986-1988

Years	1986	1987	1988
Volumes Assigned	288	314	N/A
Datasets	10,610	10,651	8,318
Space Allocated (Megabytes)	19,152	20,480	17,752
Space Used (Megabytes)	13,123	14,654	15,276
Utilization	69%	72%	86%

In 1986, 288 volumes contained 10,610 data sets with 19,152 megabytes allocated. Of that space, 13,123 megabytes were used. This is 69% of the space allocated. At that time, 66% of the available space was allocated and only 45% of the available space was used.

In 1987, 314 volumes contained 10,651 data sets with 20,480 megabytes allocated. Of that space, 14,654 megabytes were used. This is 72% of the space allocated. At that time, 65% of the available space was allocated and only 47% of the available space was used.

In 1988, under DFHSM, there were 6,754 data sets migrated on 3480 tape volumes and 1,564 data sets online on 3380 DASD for a total of 8,318. This is a 22% decrease in the number of data sets from 1987. This might be explained by the review and deletion of obsolete data which took place prior to the migration off the MSS. There were 10,777 megabytes in the data sets which were migrated and 4,499 megabytes in data sets online for a total of 15,276 megabytes in use, which is a 4% increase over 1987. Space used is 86% of the space allocated. Of the 6,975 megabytes allocated for online data sets, 64.5% was used.

Table 4 shows the results of an evaluation conducted in the first week of May, 1986. This evaluation was not re-run in 1987. From Table 3, it would appear to be unnecessary to re-run it because the data for 1987 were quite similar to 1986. This information was used as a starting point to determine how much space would be needed for primary volumes under DFHSM. This table shows that approximately 24% of the available space was referenced within 31 days. Only 15% of the available space was referenced within seven days. This implied an initial value of 6.99 gigabytes of space needed on primary volumes under DFHSM for data sets to be used in less than 31 days.

Table 4. SPACE UTILIZATION BY DAYS SINCE LAST REFERENCED—MAY, 1986

DAYS SINCE REFERENCED	Gigabytes Used	% Total Space Allocated
0	1.456	5
0-2	1.747	6
3-5	.582	2
6-7	.582	2
8-15	1.165	4
16-31	1.456	5
32-90	2.912	10
91-365	4.368	15
Over 365	4.951	17
Total	19.219	66

Added to this space requirement was space needed for all the temporary data sets on the system. At the time this evaluation was made, there were six 3350 volumes (or 1.8 gigabytes) dedicated to this purpose. Initially (June 1987), ten 3380 volumes were

allocated as DFHSM primary volumes for a total of 12.6 gigabytes of space. By the end of July, three migration level 1 volumes (3.78 gigabytes) were added. Migration level 2 is on 3480 tapes. Fourteen months after the migration from the MSS, seventy-six (76) 3480 tapes are being used. Two hundred were estimated to be needed for saving data up to two years after last reference. This estimate still is reasonable. Fourteen months after the initial assignments, one 3380 volume has been removed from the primary volumes to free it for other use. This caused much interval migration to occur on the other primary volumes, initially. Interval migration occurs when a primary volume reaches 80% of full capacity. DFHSM checks each volume, hourly, to see if interval migration is needed. The migration parameter of ten days since last reference on the primary volume has been changed to six, and further modifications may be necessary. Evaluation will continue and changes will be made when the need arises.

B. USAGE

Table 5 shows the percentage of the total amount of space used in data sets of various sizes for each of the three years. In 1986 and 1987, most of the data sets were quite small because of the limits of the MSS. The 1988 data indicate a good distribution through the range of sizes. This shows that, after the migration, the user was able to decide the optimum size data set for the particular application.

Table 5. PERCENTAGE OF DATA SETS OF VARIOUS SIZES FOR 1986 TO 1988

DATA SET SIZES	1986	1987	1988
0-10	67.0	65.0	31.0
11-20	15.0	19.0	16.0
21-30	8.0	8.0	7.0
31-60	3.0	6.0	11.0
61-100	6.0	1.0	7.5
101-150	0.3	1.0	6.5
151-200	0.1	0.1	6.0
201-250	0.4	0.3	3.0
251-300	0.0	0.0	1.0
Over 300	0.0	0.0	11.0
Data set sizes are in 3380 tracks. Table entries are the percentages of allocated space.			

In June 1988, a study was made of the use of disk storage by the large MVS users, i.e., those using more than 25 CPU hours during the winter quarter, January to March, 1988. The results are presented below in a series of tables based on CPU utilization. In each case, the table entries are the total numbers of 3380 tracks occupied by the user's data sets of specified sizes. Only data sets labelled with the user's id number were counted. Users were not interviewed to determine if they used other data sets. A total of 17 users were involved in this study.

In Table 6, student user S3242 (Air Ocean Sciences) used over 400 CPU hours on MVS for the first quarter 1988. He graduated in June 1988 and primarily used large files for his data sets. Faculty user, F3862 (Assistant Professor, Oceanography) used between 301 and 400 CPU hours in the first quarter. She used a wider range of sizes but, also, primarily used large files.

Table 6. DISTRIBUTION OF DATA SETS FOR USERS WITH OVER 300 CPU HOURS

DATA SET SIZES	USER	
	S3242	F3862
0-10	0	0
11-20	0	0
21-30	0	25
31-60	0	0
61-100	0	0
101-150	0	268
151-200	534	193
201-250	0	686
251-300	0	0
Over 300	15,917	12,848
TOTAL	16,451	14,020
Data set sizes and user amounts		
are in 3380 tracks.		

In Table 7, faculty users, F4073 (Visiting Professor, Meteorology) and F3922 (Adjunct Professor, Physics) and student users, S3048 (June 1988 graduate) and S3085 (September 1988 graduate) both in Air Ocean Sciences, used between 101 and 200 CPU hours in the first quarter. As can be seen in Tables 7 and 8, the amount of space and size of data sets varied greatly.

Table 7. DISTRIBUTION OF DATA SETS FOR USERS WITH 101-200 CPU HOURS

DATA SET SIZES	USER			
	F4073	S3085	F3922	S3048
0-10	0	0	61	0
11-20	0	0	182	0
21-30	0	0	99	132
31-60	120	0	344	0
61-100	2,097	0	90	0
101-150	8,707	0	0	0
151-200	4,044	0	0	0
201-250	0	0	0	0
251-300	0	0	0	0
Over 300	540	5,027	0	0
TOTAL	15,508	5,027	776	132
Data set sizes and user amounts are in 3380 tracks.				

The users in Table 8 consumed between 51 and 100 CPU hours on MVS during the first quarter 1988: faculty users, F3902 (Meteorologist) and F1950 (Associate Professor, Mechanical Engineering) and student users, S4555 and S1709, both in Naval Engineering.

Table 8. DISTRIBUTION OF DATA SETS FOR USERS WITH 51-100 CPU HOURS

DATA SET SIZES	USER			
	F3902	S4555	F1950	S1709
0-10	268	4	1	8
11-20	2,473	0	20	56
21-30	1,849	0	30	29
31-60	193	135	90	108
61-100	9,719	0	0	246
101-150	1,770	110	101	2,640
151-200	2,476	0	159	0
201-250	893	230	0	0
251-300	0	0	0	0
Over 300	2,508	0	0	2,500
TOTAL	22,149	479	401	5,587
Data set sizes and user amounts are in 3380 tracks.				

The users in Table 9 used 25-50 CPU hours. They include faculty users, F3956, (NAVSEA Professor, Meteorology), F3971, (Assistant Professor, Oceanography), F0455, (Professor, Physics), and F2074, (civilian staff, Meteorology), NPS staff user, N3958, (Program Manager, PERSEREC), and student users, S4550, (Naval Engineering), and S3064, (Operational Oceanography), both June 1988 graduates.

Table 9. DISTRIBUTION OF DATA SETS FOR USERS WITH 25-50 CPU HOURS

DATA SET SIZES	USER						
	F3956	S4550	F3971	F0455	N3958	F2074	S3064
0-10	8	9	11	9	184	0	0
11-20	45	0	535	19	231	0	15
21-30	0	0	354	90	97	0	60
31-60	131	0	231	99	60	0	0
61-100	100	0	1,672	836	267	1,080	0
101-150	127	0	0	369	284	1,239	150
151-200	152	0	4,408	912	0	0	0
201-250	228	0	1,140	684	872	4,316	0
251-300	0	300	0	285	293	525	0
Over 300	5,900	0	1,596	12,147	1,271	5,276	532
TOTAL	6,691	309	9,947	15,450	3,559	12,436	757
Data set sizes and user amounts are in 3380 tracks.							

APPENDIX B. ACRONYMS

APAR	Authorized Program Analysis Report of IBM program errors filed by users
Cache	High-speed buffer storage that contains frequently accessed instructions and data, usually on solid-state components
CPU	Central Processing Unit
DASD	Direct access storage device, a device on which access time is effectively independent of the location of the data
DFDSS	Data Facility Data Set Services, a DASD data and space management tool
DFHSM	Data Facility Hierarchical Storage Manager, an IBM program product (number 5665-329) that uses space management, backup, and recovery to manage data sets on a hierarchy of storage devices
GUIDE	IBM users' group for DOS operating systems
IBM	International Business Machines
IBM 370/145	CPU introduced in 1970
IBM 370/168	CPU introduced in 1972
IBM 2321 Data Cell	A direct access storage volume containing strips of tape on which data are stored [Ref. 18].
IBM 3084	CPU introduced in 1982
IBM 3090	CPU introduced in 1985
IBM 3350 DASD	317.5 megabytes capacity with 1.198 megabytes per second transfer rate, uses a sealed head/disk assembly with 16 recording surfaces
IBM 3380 DASD	2.5 gigabyte capacity with an average seek time of 16 milliseconds and a data transfer rate of 3 megabytes per second; a film head technology is used to achieve writing and reading of data recorded at higher densities than previous disk storage devices
IBM 3480	Cartridge tape drive subsystem which consists of a buffered, microprocessor-controlled, control unit and two microprocessor-controlled, tape drives that use a cartridge-enclosed 18-track, high-density magnetic tape cartridge; data rate of 3.0 megabytes-per-second.
IBM 3850	Hardware for the Mass Storage Subsystem
IBM 3880	High-performance cached DASD subsystem; the model 11 is used for paging and swapping and can be attached to a

	1.5, 2.0, or 3.0 megabytes per second channel and to one string of two, four, six, or eight 3350 devices; the model 13 is designed for system and application data and can be attached to 3.0 megabytes per second data-streaming channels and 3380 DASD.
ICF	Integrated Catalog Facility, offers significant advantages over, and designed as a functional replacement for, OS control volumes (CVOLs) and VSAM catalogs.
JCL	Job control language used to identify a job to an operating system and to describe the job's requirements
ML1	Migration Level 1, category of volume to which DFHSM migrates data sets from primary volumes
ML2	Migration Level 2, category of volume to which DFHSM migrates data sets from migration level 1 or primary volumes
MSS	Mass Storage Subsystem, composed of an IBM 3850 and supporting staging disks, IBM 3350s
MVS	Multiple Virtual Storage; IBM batch operating system
NPS	Naval Postgraduate School
Primary volume	Category of DFHSM volume containing data sets that are directly accessible to the users and managed by DFHSM
SHARE	Users' group for users of large IBM systems
TSO	Time Sharing Option for interactive use under MVS
VM	Virtual Machine (interactive operating system)
VSAM	Virtual Storage Access Method for indexed or sequential processing of fixed and variable-length records on direct access devices. Files may be organized in a logical sequence by means of a key field or a relative-record number.

LIST OF REFERENCES

1. IBM Internal Use Only Memo of February 12, 1985, Information Systems Group, International Business Machines Corporation. "IBM 3090 Processor Complex" and "3850 Mass Storage System Information," Rye Brook, New York: 1985.

2. Moore, Fred, "Improvements in Architecture and Technology are Resulting in a New Storage Hierarchy," *Mainframe Journal*, pp. 15-26, 64-65, Dallas, Texas, September-October 1987.

3. International Business Machines Corporation, *Introduction to IBM 3880 Storage Control Model 13, GA32-0062-1*, 2nd Ed., IBM Corporation, 1983.

4. Terkowitz, Roberta S., "Different Storage Technologies Meet Different Needs," *Government Computer News*, March 27, 1987.

5. Information Systems Group, International Business Machines Corporation. IBM Product Announcement, "IBM 3990 Storage Control," Rye Brook, New York: 1987.

6. Profit, L.A., "DASD Subsystem Management," *SHARE 67 Proceedings*, pp. 480-492, Chicago, Illinois, August 1986.

7. Goldstein, Stephen, "Storage Performance: An Eight Year Outlook," IBM speaker at Computer Measurement Group meeting in San Francisco, California, September 15, 1987.

8. International Business Machines Corporation, *Introduction to the IBM 3850 Mass Storage System (MSS)*, 3rd Ed., IBM Corporation, 1975.

9. International Business Machines Corporation, *IBM 3380 Direct Access Storage: General Information, GC26-4193-2*, 3rd Ed., p. 11, IBM Corporation, 1986.

10. International Business Machines Corporation, *IBM 3350 Direct Access Storage: Description Manual, GA26-1638-4*, 5th Ed., p. 17, IBM Corporation, 1984.

11. Malarkey, D. J., and Taylor, L. B., *3850 Mass Storage Subsystem Migration Planning, GG66-0208-0,*, Washington Systems Center, IBM Corporation, Gaithersburg, Maryland, 1985.

12. International Business Machines Corporation, *OS/VS Mass Storage System Extensions Services: Guide, SH35-0035-1,*, 2nd Ed., IBM Corporation, 1982.

13. International Business Machines Corporation, *Operator's Library: IBM 3850 Mass Storage System (MSS) Under OS/VS, GC35-0015-4*, IBM Corporation, 1983.

14. International Business Machines Corporation, *Reference Manual for IBM 3850 Storage Control Model 1 and IBM 3330 Disk Storage, GA26-1592-5*, 6th Ed., IBM Corporation, 1976.

15. International Business Machines Corporation, *IBM 3380 Direct Access Storage: Planning and Use, GC26-4208-0*, 1st Ed., p. 146, IBM Corporation, 1985.

16. International Business Machines Corporation, *MVS/370 Catalog Administration Guide, GC26-4053-1*, 2nd Ed., p. 30, IBM Corporation, 1985.

17. International Business Machines Corporation, *Data Facility Hierarchical Storage Manager Installation and Customization Guide, SH35-0084-1*, 2nd Ed., p. 23, IBM Corporation, 1985.

18. International Business Machines Corporation, *Dictionary of Computing, SC20-1699-7*, 8th Ed., IBM Corporation, 1987.

INITIAL DISTRIBUTION LIST

		No. Copies
1.	Defense Technical Information Center Cameron Station Alexandria, VA 22304-6145	2
2.	Library, Code 0142 Naval Postgraduate School Monterey, CA 93943-5002	2
3.	Linda Mauck Code 0141 Computer Center Navy Postgraduate School Monterey, CA 93943	2
4.	W.R.Church Computer Center Code 0141 Navy Postgraduate School Monterey, CA 93943	1
5.	Commanding Officer Naval Regional Data Center Washington Navy Yard Washington, D. C. 20374-1662	1
6.	Technical Director Naval Data Automation Command Washington Navy Yard Washington, D. C. 20374-1662	1